Paepcke, S. & Takayama, L. (2010). Judging a bot by its cover: An experiment on expectation setting for personal robots. Proceedings of Human-Robot Interaction Conference: HRI 2010, Osaka, JP, 45-52.

Judging a Bot By Its Cover:

An Experiment on Expectation Setting for Personal Robots

Steffi Paepcke & Leila Takayama
Willow Garage
68 Willow Road, Menlo Park, CA USA
{steffi, takayama}@willowgarage.com

***Abstract*—Managing user expectations of personal robots becomes particularly challenging when the end-user just wants to know what the robot can do, and neither understands nor cares about its technical specifications. In describing what a robot can do to such an end-user, we explored the questions of (a) whether or not such users would respond to expectation setting about personal robots and, if so, (b) how such expectation setting would influence human-robot interactions and people's perceptions of the robots. Using a 2 (expectation setting: high vs. low) x 2 (robot type: Pleo vs. AIBO) between-participants experiment (*N*=24), we examined these questions. We found that people's initial beliefs about the robot's capabilities are indeed influenced by expectation setting tactics. Contrary to the hypotheses predicted by the Self-Fulfilling Prophecy and Confirmation Bias, we found that erring on the side of setting expectations lower rather than higher led to less disappointment and more positive appraisals of the robot's competence.**

Keywords-human-robot interaction; user expectations

I. Introduction

First impressions can have a lasting effect on a user's future interactions with a product, and this may hold especially true for such novel products as robots. Because everyday interaction with robots is not commonplace, our perceptions and expectations of what a robot can and might do are shaped, at least in part, by what we see in movies such as the Terminator series or WALL-E. Robot users may be led astray by these preconceived notions set by Hollywood, resulting in disappointing or perhaps, more ideally, satisfying human-robot interactions. From a marketing standpoint, the former outcome is suboptimal, and every effort should be made to deliver a satisfying user experience. First impressions often affect one's final appraisal of a person or system [1], making expectation management important for facilitating positive interactions and continued use. In the interest of being as honest as possible about a robot's capabilities, it is tempting to provide an exhaustive list of technical specifications for the robot.

An enumeration of technical specs of the Pleo robot indicate that it has IR Communication, Infrared Interruptor, Chin Touch Sensor, Front Speaker, NiMH Rechargeable Battery, Tilt and Shake Sensors, Ground Sensors, Leg Touch Sensors, Force Feedback Sensor, Rear Touch Sensor, Shoulder Touch Sensor, Binaural Microphones, Color Camera, Head Touch Sensor, and Rear Speaker [2].

Similarly, the Sony AIBO ERS-220 is advertised as "an entertainment robot able to walk on four legs. It has a total of 16 actuators throughout its body to control its movements, and 19 lights on its head, tail, and elsewhere to express emotions like happiness or anger and reactions to its environment" [3]. While listing technical specifications and the number of actuators is informative to robotics hobbyists, it is less informative to end-users, who just want to know what the robot is for and what it can do.

How is this best accomplished? If a robot is introduced as less capable than it in fact is, users may be pleasantly surprised when it exceeds their expectations. This would be consistent with American businessman and writer Tom Peters' formula for success: under-promise and over-deliver. According to Peters, "customers unfailingly prefer slightly less aggressive promises… that are honored" [4]. The idea is as follows: companies that make modest promises but go on to surpass these promises, will experience higher customer satisfaction than companies that do not set their customers' expectations slightly lower from the outset. These promises include project completion dates, budget estimates and customer service wait times. If a customer's expectations are exceeded, they will, according to Peters, keep coming back with their business. This is a well-known "management maxim" seen in text books (e.g., [5, 6]) that is repeated by companies such as AOL and Southwest Airlines, though it also has its opponents (e.g., [7]).

In contrast, if a robot is introduced as more capable than it in fact is, users could be influenced by a confirmation bias [8] that makes setting high expectations a better strategy.

The risk of under-selling the robot is that the self-fulfilling prophecy [9] could make the user truly believe that the robot is worse than it actually is. The risk of over-selling the robot is that the user might become disappointed when the robot does not live up to its advertised capabilities.

Actual marketing descriptions of these robots seem to err on the side of over-selling the robots:

> "Developed to encourage human and robot interaction, AIBO introduces you to new pleasures and lifestyles... In Japanese, the word 'aibou' means 'partner' or 'pal'." [3]

> "Like any creature, Pleo feels hunger and fatigue—offset by powerful urges to explore and be nurtured. He'll graze, nap and toddle about on his own—when he feels like it! Pleo dinosaur can change his mind and his mood, just as you do." [10]

While it makes sense to entice a potential buyer into purchasing one of these robots, there is also the risk that consumers will become disappointed if the robot cannot live up to the promises made by the advertising campaign.

With so many possible outcomes from both under- and over-selling these robots, the easiest solution is to just set the expectations honestly and "right." Unfortunately, this is quite difficult and may not be as successful as intended. There is no simple, standard measure that can be used to advertise how well a robot can sense touch, navigate while avoiding obstacles, or interact with people. A detailed technical description is required in order to accurately convey this information, but by and large, non-engineer end-users may not find this information engaging or useful. Additionally, end-users will not possess identical technical backgrounds, so a list of technical specifications will yield varying expectations across different users. Thus, a straightforward index of robot components will likely not produce an ideal introduction to a robot for the majority of novice users.

II. Related Concepts

Several concepts from the social sciences informed our inquiry into expectation setting for robots. From sociology, we use the concepts of face and face work; from social psychology, we use the concept of the self-fulfilling prophecy; from cognitive and social psychology, we use the concept of the confirmation bias.

A. *Face and face work*

The notion of "face" was largely explicated by Erving Goffman [11], who wrote, "The term face may be defined as the positive social value a person effectively claims for himself by the line others assume he has taken during a particular contact" (p. 5). Face is a presentation of self [12] that one presents to a particular audience in a particular setting. The face you present to your parents is typically different from the face you present to social acquaintances or co-workers; none of those faces are false, but they are different because "face" is socially defined in interactions between people.

Face work involves "the actions taken by a person to make whatever he is doing consistent with face" (p. 12); those who are said to have tact, savoir-faire, and diplomatic skills are people who engage in very effective face work (p. 13) [11]. Face work could be considered as a form of impression management [13]. There is a side of ourselves that we want to present to each of the different social groups we belong to, which is why some social psychologists even conceptualize the "self" as being an interpersonal creation rather than a monolithic entity [14].

These concepts of face and face work lay the groundwork for thinking about how robots could and should present themselves (or be presented) to different audiences. Because robotics hobbyists, developers, and researchers are interested in knowing what kinds of technical specifications a robot has (e.g., the types of sensors, actuators, joint limits), it is appropriate for the robot's face toward that audience to be framed as such. However, because consumers and other end-users of personal robots are more likely to be interested in what the robot can actually do, it is more appropriate for the robot to be presented in terms of its capabilities and uses.

B. *Self-fulfilling Prophecy*

The self-fulfilling prophecy describes how social beliefs (e.g., high or low expectations for a student's intelligence) influence how we interact with one another (e.g., teachers calling more on the "gifted" students), and sometimes even how others ultimately perform (e.g., "gifted" students performing better) [9]. Students who held high expectations of their teacher's competence after hearing that, "Professor Smith is interesting," found the teacher to be more interesting and even learned more from those teachers than students who heard that, "Professor Smith is a bore" [15]. Simply manipulating what one hears about another person influences both perceptions and behavioral outcomes in interacting with others.

Given that so many rules from social psychology apply to the interactions between people and interactive media [16, 17], it is quite possible that the expectations we bring to interacting with robots will influence our final perceptions of those robots and the ways in which we interact with them. This presents us with a self-fulfilling prophecy hypothesis (H2), which predicts that users will spend more time interacting with robots when expectations of the robots are high.

C. *Confirmation Bias*

From the cognitive psychology literature, there is a phenomena related to the self-fulfilling prophecy called the confirmation bias. The confirmation bias is also known as the confirmatory bias. Both terms refer to the tendency for people to seek or interpret evidence in ways that favor one's existing beliefs or expectations [8]. This is a known bias that could explain how people become polarized on opposite sides of debates. If you show the same piece of evidence to proponents and opponents of complex social issues such as the death penalty, both sides view the evidence that supports their own beliefs as more convincing, whereas evidence that competes against their own beliefs is judged more critically [18].

In light of what we know about how people tend to seek evidence that confirms the beliefs they already hold, we hypothesize that people might also seek evidence that confirms the beliefs they already hold about a robot's capabilities. Indeed, higher expectations have been found to predict better consumer perceptions of service quality [19]. If a person is over-sold on a robot (e.g., told that a robot has many features that enable it to interact with people), then that person might be more inclined to perceive more competence when interpreting the robot's behaviors. In the current study, it was not uncommon for people to say a command, see the robot begin a new behavior, and express delight that the robot obeyed the command; in many cases, the robot was not actually responding to the command at all, but had merely, by happenstance, begun some new action routine at that moment. This brings us to the hypothesis (H3A) that people will perceive the robots more positively when their expectations of the robot are set high rather than low. This confirmation bias hypothesis contrasts against the formula for success of under-promising and over-delivering, which is its competing hypothesis (H3B).

III. Research Question and Hypotheses

The design questions at hand are: Do end-users notice when robot capabilities are explained to them? If so, do those descriptions influence end-user perceptions and behaviors with the robots? Is it better to over-promise or under-promise when setting expectations about interactive robots?

The research questions at hand are: Does expectation setting matter in human-robot interaction? If so, how does expectation setting around robots influence human perceptions and interactions with these robots? Do they follow in the steps of psychological theories of the self-fulfilling prophecy and confirmation biases, or do they follow in the steps of the business philosophy of under-promising and over-delivering?

Based upon theories from sociology, cognitive psychology, social psychology, and business, we embarked on this project with the following research hypotheses:

(H1) Expectation Setting: Expectation setting will influence human-robot interaction, i.e., people's beliefs will be influenced by expectation setting strategies used in presenting the robots

(H2) Self-fulfilling Prophecy: Setting user expectations of the robot's capabilities high will engage the user more in interacting with the robot than when they are set low, i.e., people will interact longer with the robots when expectations are set high rather than low

(H3A) Confirmation Bias: Setting user expectations of the robot's capabilities high will create more positive perceptions of the robot than when they are set low, i.e., people will view the robots more positively when expectations are set high rather than low

(H3B) Under-promise and Over-deliver: Setting expectations of the robot's capabilities low will create more positive perceptions of the robot than when they are set high, i.e., people will view the robots more positively when expectations are set low rather than high

IV. Study Design and Methods

In a 2 (robot type: Pleo vs. AIBO) x 2 (expectation setting: high vs. low) between-participants experiment (N=24), we studied the effects of robot type and user expectation on system appraisal and user experience. Table 1 presents the distribution of participants in the study's design.

TABLE I. User Expectations of Robots Experiment Design

Between-participants Variables		**Expectation Setting**	
		High	***Low***
Robot Type	***AIBO***	3 women and 3 men	3 women and 3 men
	Pleo	3 women and 3 men	3 women and 3 men

Because we were primarily interested in how expectation setting with personal robots influences human-robot interactions, we chose to run 12 participants in each of those two conditions. Other variables such as robot type (AIBO and Pleo) and participant gender (male and female) were balanced across those conditions.

A. Participants

Twenty-four volunteers participated in this experiment (12 male and 12 female, balanced across conditions) and each received a $10 gift card. Participants were recruited via local mailing lists. Participants' ages ranged from 20 to 60 years of age (M=30.46, SD=12.07). Their occupations included student, human resources administrator, IT specialist, consumer marketing assistant, and non-profit recruiter. Their personalities were generally characterized as follows: Extraversion (M=1.43, SD=3.15), Agreeableness (M=2.50, SD=2.23), Conscientiousness (M=3.26, SD=2.25), Emotional stability (M=2.46, SD=2.04), and Openness to Experience (M=4.25, SD=1.42) [20, 21]. None of the participants had ever before interacted with the type of robot they encountered in the study.

B. Experiment Manipulation

While the study design involved balancing the conditions with each type of robot, Pleo and AIBO, the real experiment manipulation of interest was user expectation settings. Participants in the condition involving high expectation setting were informed by the experimenter that, "This robot has many people-sensing and interactive capabilities." Participants in the condition involving low expectation setting were informed by the experimenter that, "This robot does not have many people-sensing and interactive capabilities." To ensure that participants would not forget this, a sign was left in the room during the participant's interaction with the robot that reminded them to think aloud (i.e., outwardly verbalize the stream of thoughts that are going through their heads as they come to mind) and remind them about the robot's supposed capabilities. The sign read either, "This robot has many people-sensing and interactive capabilities," or, "This robot does not have many people-sensing and interactive capabilities."

C. Materials and Measures

To improve the generalizability of this study, we used two types of robots instead of only one. We chose to use a Sony AIBO ERS-220 with an AIBO Explorer memory card, and an Ugobe Pleo. Both robots were designed to interact with people, are relatively similar in size and form, and are designed to be similar to canine pets. Both are programmable and are primarily intended as robotic pet companions [22]. We did not program either robot to display any special behaviors beyond those already provided out-of-the-box; instead, we had participants freely interact with Pleo or AIBO in their default settings. See Figure 1.

To record the participants' interactions with the robots, we used a digital video camera pointed at the tabletop. This way, we were able to record the participants' voices, the robot's actions, and the amount of time people spent with the robots, M=12 minutes and 7 seconds, SD=7 minutes and 44 seconds.

Pen and paper questionnaires were administered both immediately before and immediately after each participant interacted with the robot. In these questionnaires, we used a series of identical 5-point Likert scale questions to assess participants' perceptions of the robots before and after interacting with the robot. Volunteers were asked to mark on a scale of "Definitely cannot" to "Definitely can" whether they believed the robot could, for example, sense touch, navigate, perceive people, and avoid obstacles. In our post-test questionnaire, we also used the Ten-Item Personality Inventory [21], an abbreviated variation of the Big Five [20], to investigate the possibility of personality traits affecting participant experience. Additionally, we used the Source Credibility Scale [23] to look specifically at participants' attitudes towards, and appraisal of, the robot. Demographic questions regarding age, gender, educational levels, pet ownership experience, etc., were also included.

D. Procedure

When participants arrived, they were given a consent form and liability waiver agreement to sign. If they chose to continue with the study, we provided the participant with a brief overview of the study protocol and asked each subject if they had ever interacted with an AIBO or Pleo before. It was uncommon for participants to have prior exposure to either, but if they had interacted with one before, we assigned them to the opposite robots.

After presenting AIBO or Pleo, we either told participants that the robot did, or did not, have many people-sensing and interactive capabilities. Next, a pre-test questionnaire about their first impressions of the robot was administered.

Participants were then instructed to think aloud about what they thought the robot was doing. Participants were told that the goal of the study was to learn how to design more human-readable robot behaviors, and that through the think-aloud protocol, we hoped to learn which robot behaviors conveyed meaning and which ones did not.

We instructed the volunteers to turn the robot on once we left the room, and spend as much time as they liked interacting with the robot while remembering to think aloud about their interpretations.

Once the participants were finished, we administered a paper-based post-test questionnaire.

V. Data Analysis

We used analysis of variance (ANOVA) to investigate whether setting user expectations influenced user behaviors and attitudes. In each ANOVA, we used the user expectation setting (high vs. low) as the main independent variable.

A. Independent Variables

Two independent variables were used in these analyses: user expectation setting (high vs. low) and robot type (AIBO vs. Pleo). Although the research questions were not particularly focused upon robot type, we included this independent variable in the data analyses because it was likely that the variance influenced by robot type would create noise in our data.

B. Dependent Variables

The manipulation check dependent variables were single-item measures of how much the participant believed that the robot could sense touch and how much the participant believed that the robot could perceive people. We had conducted a principal components analysis (PCA) on a set of these robot capability beliefs, but they did not cluster into a single index so we chose to analyze the items separately.

Two types of dependent variable indices were calculated from the raw data in order to conduct these analyses.

The first index was an ordinal value (-1, 0, or +1) of whether a person's expectations of the robot decreased, remained the same, or increased between the time when the person first encountered the robot and the time after the person decided to stop interacting with the robot. This was calculated for perceptions of both people-perceiving and touch-sensing abilities.

The second index was calculated for perceived competence of the robot, which was an average of how experienced and informed the robot seemed to be; this was done because two items of "experienced" and "uninformed" (reverse coded) from the Source Credibility Scale were correlated (r=-.62, p<.01).

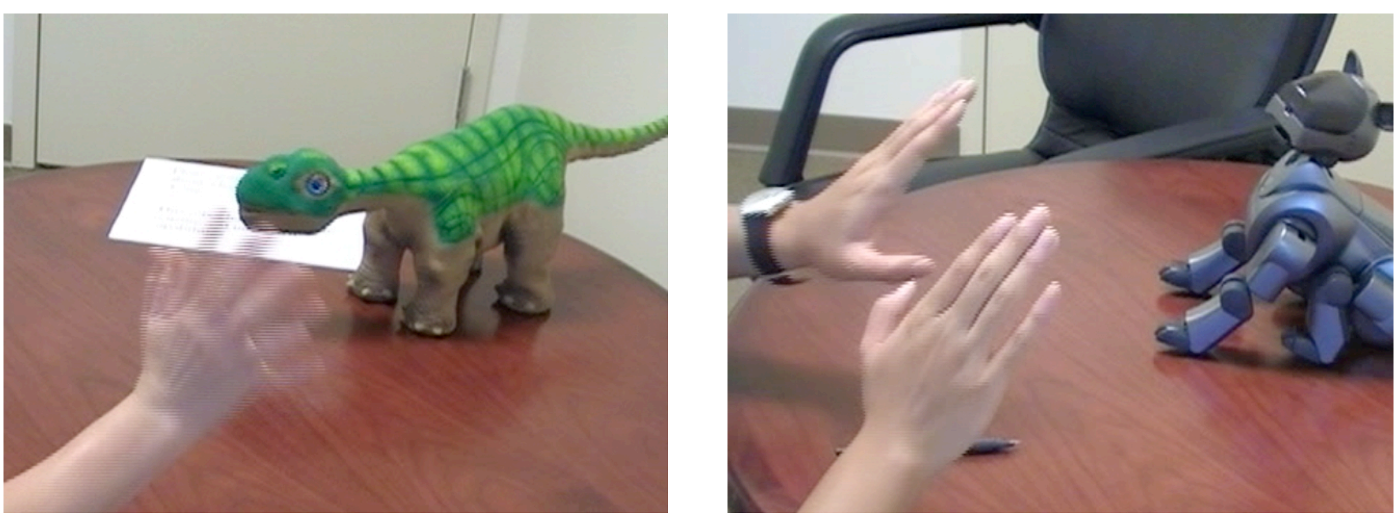

(1a) (1b)

Figure 1. (a) Pleo and (b) AIBO robots in the experiment setting; participants often waved their hands around to see if the robot could sense them

Finally, we also measured how many seconds each person interacted with the robot, starting the timer from the moment that the experimenter left the room to the moment when the participant stood up to tell the experimenter that she or he was done interacting with the robot.

C. Results

1) Manipulation Checks

Our manipulation of setting people's expectations high vs. low influenced participants' perceptions of how capable the robot would be at perceiving people and sensing touch. Before interacting with the robots, participants whose expectations were set high believed the robot would be better at perceiving people (M=2.83, SE=0.24) than participants whose expectations were set low (M=1.75, SE=0.25), $F(1,22)=9.73$, $p<.01$. See Figure 2a. Before interacting with the robots, participants whose expectations were set high also believed that the robot would be better at sensing touch (M=4.25, SE=0.22) than people whose expectations were set low (M=3.33, SE=0.36), $F(1,22)=4.84$, $p<.05$. See Figure 2b.

2) Effects of Expectation Setting

Expectation setting affected a person's perceived competence of the robot as well as the amount that a participant's expectations of the robot changed over time.

After interacting with the robots, participants whose expectations were set high experienced a decrease in perceived touch-sensing capabilities of the robot (M=-0.42, SE=0.22), whereas participants in the low expectation condition experienced an increase in perceived touch-sensing capabilities (M=0.25, SE=0.36), $F(2,21)=2.12$, $p<.05$. See Figure 2c.

Robot type (AIBO vs. Pleo) also influenced how much the expectation of touch-sensing capabilities of the robot changed before vs. after interacting with the robot, $F(2,21)=2.12$, $p<.05$; AIBO (M=0.58, SE=0.24) vs. Pleo (M=0.67, SE=0.27).

The differences caused by expectation settings were not found to significantly affect changes (before vs. after interacting with the robot) in apparent people-perceiving capabilities, $F(2,21)=0.49$, p=.63.

Perceived competence of the robot was also influenced by expectation setting. After interacting with the robots,

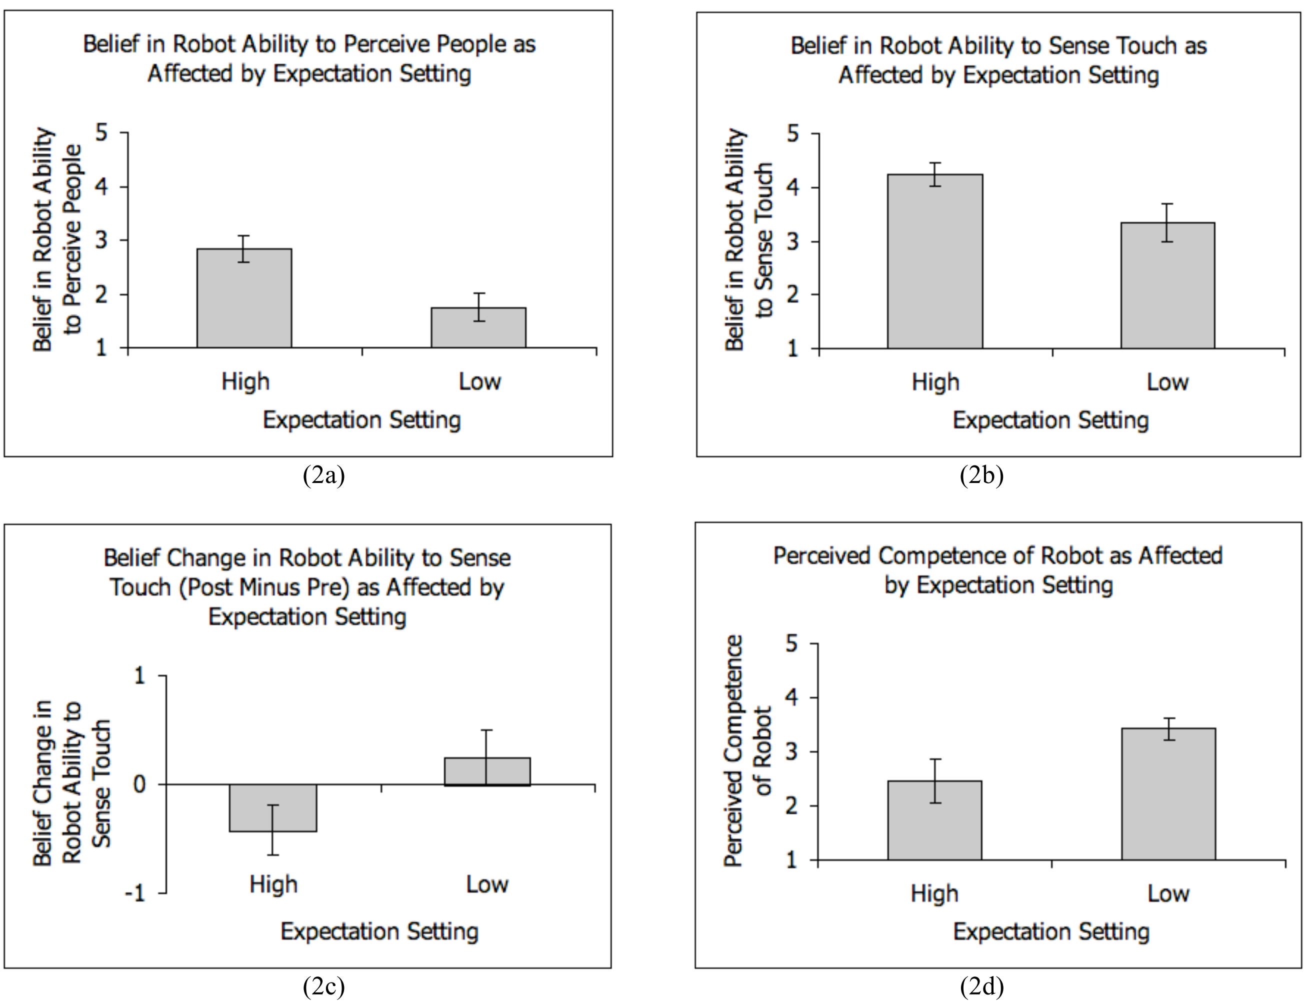

Figure 2. These bar charts represent the influences of expectation setting upon pre-interaction expectations of the robot's abilities to (2a) sense touch and (2b) perceive people; they also represent the influence of expectation setting upon post-interaction (2c) changes in beliefs about the robots' abilities to sense touch (negative values indicate disappointment) and (2d) final perceived competence of the robot; mean and standard error values

participants whose expectations were set high believed that the robot was less competent (M=2.46, SE=0.41) than participants whose expectations were set low (M=3.42, SE=0.21), $F(1,22)=2.09$, $p<.05$. See Figure 2d.

Even though we did not constrain the amount of time that a person could spend with the robot, the current study did not find any significant differences of expectation setting affecting how long people interacted with the robots, $F(2,21)=0.99$, $p=.33$; robot type was not found to influence how much time people interacted with the robot either, $F(2,21)=0.72$, $p=.40$.

VI. Discussion

A. *Quantitative Results and Interpretations*

Because these robots were designed and marketed specifically for interacting with people, we believed that they would have at least some people-sensing and interactive capabilities. By setting expectations either high or low with respect to these capabilities, we aimed to see if being too humble or too boastful about the robot's capabilities would influence users' interactions with, and perceptions of, the robots. The manipulation checks for how well people believed the robot could sense touch and perceive people showed that setting expectations indeed influenced beliefs about the robots' capabilities. These findings support Hypothesis 1, which stated that people's beliefs would be influenced by expectation setting strategies used in presenting the robots.

More importantly, setting expectations of the robot's interactive capabilities too high made people feel even more disappointed in the robot's touch-sensing capabilities (before vs. after interacting with the robot) and made people ultimately perceive the robot as being less competent than when setting expectations low. These findings support Hypothesis 3B, which predicted that setting expectations of the robot's capabilities low would create more positive perceptions of the robot than when they were set high. This provides evidence against Hypothesis 3A that high expectations yield more positive perceptions of the robots.

These data analyses did not find support for Hypothesis 2, which stated that people would interact longer with the robots when expectations were set high rather than low (consistent with the self-fulfilling prophecy). We only measured the time users spent interacting with the robot; it is possible that other measures of engagement would have been more appropriate than time-on-task.

As part of the post-study questionnaire, we asked an open-ended question about whether or not the robot lived up to the participants' expectations. Among people in the high expectation setting group, we found responses such as, "No, I expected the dog (AIBO) to at least respond to the same verbal commands that a real dog would—such as come, sit, stay, etc." Among people in the low expectation setting group, we found responses such as, "I was amazing [sic] how much he interacted with me and responded to touch. I found myself talking to him like I would a dog!" While the physical form of AIBO and Pleo seemed to bring about expectations for dog-like behaviors from the robots, we still found that setting expectations about the robots' capabilities influenced the apparent capabilities of the robots and ultimate impressions of the robots for participants in this study.

B. *Qualitative Observations*

We reviewed the video data collected in this study in search of interaction patterns that might exist between the participants and the robots.

Many people in the current study seemed to expect dog-like behaviors from these personal robots. Their physical forms likely influenced this expectations because AIBO is physically designed to look like a dog and Pleo was designed to behave like one, too. Consistent with previous work [24], we found that people would often try to give commands to the robots that would be given to dogs, including: sit, stand, lie down, shake, and come. They would also pat the robot on the head and back. Indeed, pet ownership seems to influence human-robot interactions in systematic ways [25, 26]. Several participants commented in the follow-up debriefing sessions that they had found themselves interacting with the robot in ways similar to how they interact with their own pets.

Some argue that the development of these robotic pets, sometimes called "ersatz companions," is unethical in that it encourages the confusion of real social and emotional relationships with false ones [27]. Indeed, we found that these robots did elicit responses from people that looked much like pet-oriented responses. This should not be surprising when one considers the broader set of mindless social responses observed in people interacting with computers [17].

Another pattern we noticed among individual participants was that some were very hands-on in their interactions with the robots, e.g., picking them up, turning them over, patting them. (See Figure 3.) A few participants even flipped the switches and pushed buttons on the robots' undersides. However, other participants were very hands-off, never or rarely touching the robot, and preferring to say things to the robots rather than come into physical contact with them.

The most surprising observation in this study was that it was not unusual for people to read too much into the robot behaviors. For example, a person would issue a verbal command to the robot, the robot would change its behavior, and then the person would conclude (thinking aloud) that the verbal command had triggered the robot's change in behavior. This was cause for many incorrect conclusions about the abilities of the robots:

In one situation: The person asks Pleo, "Do you want this leaf?" Pleo happens to make a sound. She says, "Leaf?" and puts the leaf to Pleo's mouth. Pleo leans forward toward the leaf. She says, "Oh, that's cool. It seems to understand leaf." Pleo makes a noise. "What about tree? Cookie?"

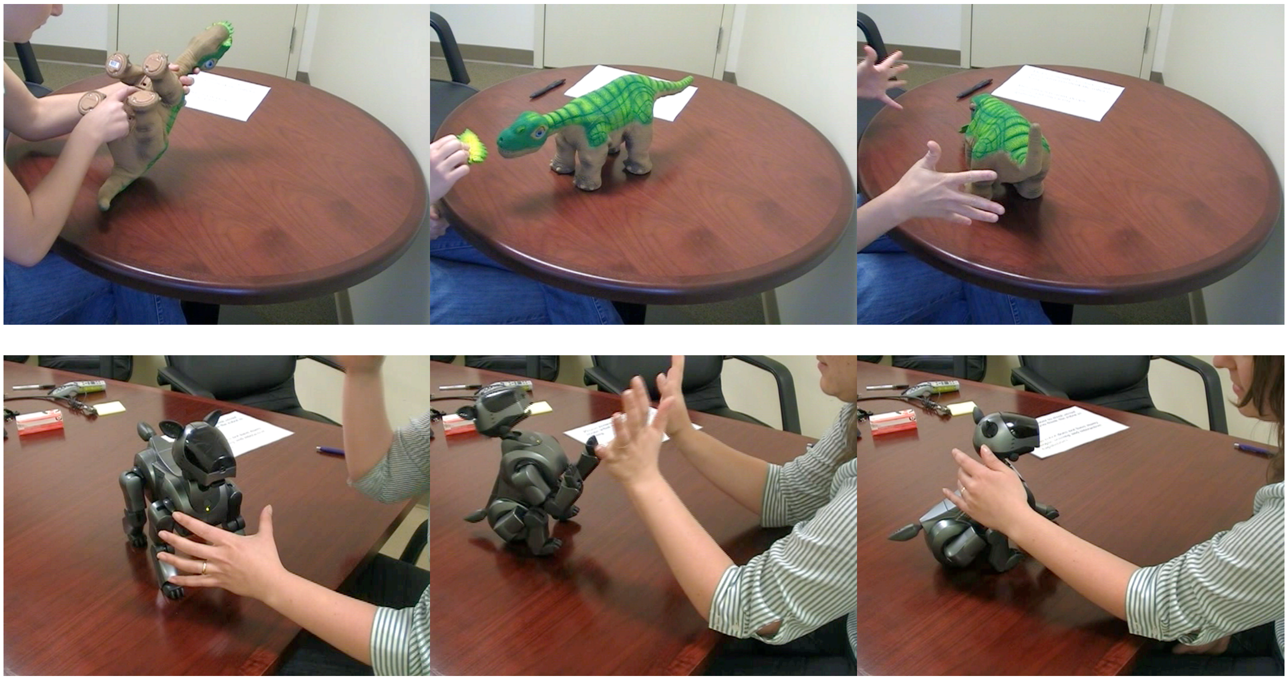

Figure 3. These sequences of interactions observed in this study depict common types of behaviors observed between participants and these robots

In yet another situation, the participant tried to turn AIBO on by pushing the power button several times. He says, "To be perfectly honest, I expected it to do something. I imagine it's turned on." He then touches the head sensor. AIBO turns on, but only because there is a time delay between hitting the power button and the robot actually moving. He says, "Oh, ok. I guess I touched this part on top and that activated it."

C. *Implications for Design*

The clearest implication for design presented in this study is that personal robots such as these should be described to end-users by erring on the side of under-promising rather than over-promising if one's goal is to mitigate disappointment in the robot's capabilities. Although people initially notice and believe descriptions of robot capabilities before they interact with the robot, on average, they become disappointed in its abilities after interacting with it and ultimately find the robot to be less competent.

D. *Limitations and Future Work*

There are several limitations to the current study that could be improved in future work. Although we aimed to generalize our results somewhat by using more than one robot, we only studied interactions with Pleo and AIBO robots. Additionally, the results found with these two robots might not extend to personal robots intended for different task domains such as assistive robotics. In this vein, another limitation is the size and animal-like appearance of the Pleo and AIBO robots. Based on findings from previous research on more or less humanoid robot heads [28, 29], it is likely that larger and/or more humanoid robots would lead participants to read different sorts of expectations into those form factors.

The current study's open-ended task of simply interacting with the robot was meant to elicit more self-motivated types of interactions from participants, but more structured, goal-oriented tasks might yield different results. We opted for the open-ended interaction because we had aimed to identify for systematic differences in how long people would engage in interactions with the robots, depending upon their expectation levels. However, goal-oriented tasks would allow one to have a clear behavioral metrics for observing and quantifying the effectiveness of the interactions. Future work could use such tasks as collaborative puzzle-solving or construction.

The particular implementation of the experiment manipulation is just one of many ways of setting expectations about robots. We chose to simply tell participants about the robot's capabilities and leave a sign on the table to remind them of its capabilities. This allowed us to be as consistent as possible with existing advertising practices of telling consumers about the product's features and also writing the list of features on the product's packaging. However, other ways to set expectations about robots include robot expressivity or providing technical specifications to users who would understand and appreciate those specifications.

Finally, future studies could investigate different user populations. The current study involved volunteers from the San Francisco Bay Area, who were not involved in robotics; we had hoped they would be somewhat representative of end-user consumers. They are at least representative of a wider age range of people than many experiments with college students. However, this population of people is not necessarily representative of broader populations (e.g., Americans, Westerners, all humans, etc.) so future work could investigate these HRI issues in other populations.

VII. CONCLUSIONS

This experiment on the effects of expectation setting with personal robots explored the question of how expectation setting would influence human-robot interactions and perceptions of personal robots. Using a controlled experiment, we found that setting expectations about a robot's capabilities indeed influenced users' beliefs about what the robot could do. However, upon interacting with the robot, people whose expectations were set high became more disappointed with the robot's capabilities than people whose expectations were set low. Furthermore, people whose expectations were set high (as opposed to low) ultimately perceived the robot as being less competent. Together, these findings suggest that people setting expectations about personal robots should err on the side of being more humble when presented to end-users.

VIII. ACKNOWLEDGMENTS

We would like to thank the volunteer participants who took the time and effort to make this study possible.

IX. REFERENCES

[1] M. Rabin and J. L. Schrag, "First impressions matter: A model of confirmatory bias," *The Quarterly Journal of Economics,* vol. 14, pp. 37-82, 1999.

[2] Ugobe, "Pleo Look Inside," in *http://www.pleoworld.com/discover/lookinside/*, 2009.

[3] Sony, "Sony manual for AIBO model ERS-220," in *http://support.sony-europe.com/aibo/1_2_library.asp*, 2009.

[4] T. Peters, "Under promise, over deliver," TPG Communications, 1987.

[5] M. Govindarajan, *Marketing Management*. New Delhi: Prentice-Hall, 2007.

[6] M. DeMoss, *The Little Red Book of Wisdom*. Nashville, TN: Thomas Nelson, 2007.

[7] R. Price and D. E. Schultz, "Promised land," *Marketing Management,* pp. 25-27, July/August 2006.

[8] R. S. Nickerson, "Confirmation bias: A ubiquitous phenomenon in many guises," *Review of General Psychology,* vol. 2, pp. 175-220, 1998.

[9] R. Rosenthal and L. Jacobson, *Pygmalion in the Classroom: Teacher Expectation and Pupils' Intellectual Development*: Crown House Publishing, 2003.

[10] Ugobe, "Pleo Discover," in *http://www.pleoworld.com/discover/discover/*, 2009.

[11] E. Goffman, *Interaction Ritual: Essays on face-to-face behavior*. New York: Doubleday, 1982.

[12] E. Goffman, *The Presentation of Self in Everyday Life*. New York: Anchor Books, 1959.

[13] B. R. Schlenker and M. F. Weigold, "Interpersonal processes involving impression regulation and management," *Annual Review of Psychology,* vol. 43, pp. 133-168, 1992.

[14] H. Markus and S. Cross, "The interpersonal self," in *Handbook of Personality: Theory and Research*. vol. xiv New York, NY: Guilford Press, 1990, pp. 576-608.

[15] D. G. Myers, *Social Psychology*, 7th ed. New York, NY: McGraw Hill, 2002.

[16] B. Reeves and C. Nass, *The Media Equation: How people treat computers, television, and new media like real people and places*. New York: Cambridge University Press/CSLI, 1996.

[17] C. Nass and Y. Moon, "Machines and mindlessness: Social responses to computers," *Journal of Social Issues,* vol. 56, pp. 81-103, 2000.

[18] C. G. Lord, L. Ross, and M. R. Lepper, "Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence," *Journal of Personality and Social Psychology,* vol. 37, pp. 2098-2109, 1979.

[19] L. O. Hamer, "A confirmation perspective on perceived service quality," *Journal of Services Marketing,* vol. 20, pp. 219-232, 2006.

[20] O. P. John and S. Srivastava, "The big-five trait taxonomy," in *Handbook of Personality: Theory and Research*. vol. 2, L. A. Pervin and O. P. John, Eds. New York: Guilford Press, 1999, pp. 102-138.

[21] S. D. Gosling, P. J. Rentfrow, and W. B. Swann, Jr., "A Very Brief Measure of the Big Five Personality Domains," *Journal of Research in Personality,* vol. 37, pp. 504-528, 2003.

[22] B. Friedman, P. H. Kahn, Jr., and J. Hagman, "Hardware companions?: What online AIBO discussion forums reveal about the human-robotic relationship," in *CHI*, Ft. Lauderdale, FL, 2003, pp. 273-280.

[23] D. Berlo, J. B. Lemert, and R. J. Mertz, "Dimensions for evaluating the acceptability of message sources," *Public Opinion Quarterly,* vol. 33, pp. 563-576, 1970.

[24] S. Kriz, G. Anderson, M. Bugajska, and J. G. Trafton, "Robot-directed speech as a means of exploring conceptualizations of robots," in *HRI*, La Jolla, CA, 2009, pp. 271-272.

[25] L. Takayama and C. Pantofaru, "Influences on proxemic behaviors in human-robot interaction," in *IROS*, St. Louis, MO, 2009.

[26] B. Mutlu, F. Yamaoka, T. Kanda, H. Ishiguro, and N. Magita, "Nonverbal leakage in robots: Communication of intentions through seemingly unintentional behavior," in *Human-Robot Interaction*, La Jolla, CA, 2009, pp. 69-76.

[27] R. Sparrow, "The march of the robot dogs," *Ethics and Information Technology,* vol. 4, pp. 305-318, 2002.

[28] C. DiSalvo, F. Gemperle, J. Forlizzi, and S. Kiesler, "All robots are not created equal: The design and perception of humanoid robot heads," in *Designing Interactive Systems*, London, England, 2002, pp. 321-326.

[29] A. Powers and S. Kiesler, "The advisor robot: Tracing people's mental model from a robot's physical attributes," in *HRI*, Salt Lake City, Utah, 2006, pp. 218-225.